

# Quantian: A single-system image scientific cluster computing environment

Dirk Eddelbuettel, Ph.D  
*Debian Project*  
edd@debian.org

Draft of 2003/12/16 21:54

## Abstract

This paper introduces the openMosix extensions to the Quantian environment for quantitative and scientific computing. Quantian, originally based on Knoppix technology, allows one to boot virtually any recent commodity i386-class computer from a single cdrom containing a compressed iso image into a fully-configured graphical workstation equipped with over 2gb of software – including 500mb of applications with a focus on quantitative or scientific computing. With the additional foundation of ClusterKnoppix providing support for openMosix-based clustering, Quantian allows one to create ad-hoc single-system image computing clusters that are already loaded with a large number of software environments commonly used in quantitatively-oriented disciplines.

## 1 Introduction

Quantian [1] is a directly bootable and self-configuring Linux system based on a compressed cdrom image. Quantian is an extension of Knoppix [2] and ClusterKnoppix [3].

The Knoppix bootable Linux system on a single cdrom (also referred to as a so-called “live cdrom”) provides a complete workstation with over two gigabytes of software along with fully automatic hardware detection and configuration.

ClusterKnoppix extends the basic Knoppix system with a kernel containing the openMosix [4] patch, user-space utilities and several additional tools. If used on a local-area network with other computers running the same kernel version and openMosix patch (but possibly of different processor architecture), an auto-discovery mechanism permits these machines to instantaneously form a single-system image computing cluster. In such a cluster, any machine can be both a client or a server. In addition, any such node can act as network-boot server to additional nodes which could be booted via PXE from the initial machine.

To this powerful computing-cluster system, Quantian

adds software with a quantitative, numerical or scientific focus: several computer-algebra systems are included, as are higher-level matrix languages, data visualization tools, and a variety of libraries and scientific applications. A particular focal point is the R [5] language and environment for which the base system as well as several add-on tools are installed.

This paper is organized as follows. In the next section, we briefly describe the single-system operation based on a Knoppix-style “live cdrom”. In the following section, we detail how the ClusterKnoppix-based automatic setup of computing clusters can transform a set of personal computers into a powerful computing cluster. An example application illustrates Quantian before a summary concludes the paper.

## 2 Single-machine use of Quantian

Quantian [1] is a directly bootable cdrom using the technology introduced by the Knoppix system [2]. While the fundamental design and setup of Quantian, and its underlying Knoppix base, has been described previously [6], we will review the key points before we outline the unique features added by Quantian.

### 2.1 The foundation provided by Knoppix

It is useful to reiterate some of the key points of the Knoppix [2] ‘live cdrom’ system:

- usable on virtually any recent desktop or laptop;
- available for system recovery and forensics due to a large number of included utilities and tools for network and system administrators;
- usable in computer labs as machines can be re-booted quickly into identical states;
- can be used by students as no system administration skills are required for setup and installation;
- lowers barriers to entry for Linux and Unix technology as it provides a ready-to-use system;
- requires little retraining as the KDE desktop environment provides a familiar user experience to other common operation system interfaces;

- provides a terminalserver mode allowing other (potentially hard-disk and cdrom-less) netboot-capable machines to be initialized and booted using the PXE protocol in a thin-client environment;
- provides a ‘persistent home’ mode where data can be written to USB storage devices (or disk partitions) to preserve states between sessions;
- permits one to try Linux risk-free as no information is written to the hard disk, or existing operating system;
- enables to try Linux on new hardware to reliably test its compatibility.

Knoppix is under active development and released several times a year.

## 2.2 Quantian contributions

The first two Quantian releases were based directly on Knoppix. In order to add quantitatively-focused software, existing programs have to be removed. The list of software that is removed from the base system is determined with the dual objectives of a) creating a maximum amount of capacity and b) removing applications with limited usefulness in a quantitative analysis setting. Consequently, we remove several large applications such as openoffice (which is also removed from the most-recent Knoppix versions), mozilla and gimp, internationalization packages, games, the bochs and wine emulators, several additional window managers, some networking tools as well as a few miscellaneous applications. In total, about 500mb of software are removed.

The base system stills contains around 1.5gb of software including a complete KDE environment with its window manager, browser, office suite, development environment and editors as well as a large number of other general-purpose tools, utilities and diagnostic applications.

Quantian then adds various sets of applications from different areas:

**mathematical** such as the giac, ginac, maxima, gap, pari and yacas computer-algebra systems, as well as the euler and xppaut applications;

**statistical** such as the R language and environment for ‘programming with data’ (along with several packages from CRAN [7], ESS (Emacs Speaks Statistics) and the Ggobi data visualization tool) as well as autoclass, gretl, mcl, multimix, x12a and xlistat;

**visualization** such as the OpenDX and Mayavi visualizers, as well as gnuplot, grace, gri, xfig and plotutils for scientific plotting;

**libraries** such as GSL (the GNU Scientific Library) and QuantLib, a quantitative finance library for risk management and trading;

**matrix environments** such as Octave, a matrix programming language and interactive environment included along with several add-on packages, scientific and numeric python, the perl data language (PDL) and yorick;

**typesetting** systems such as the lyx and kile frontends to L<sup>A</sup>T<sub>E</sub>X, as well as the auctex mode for XEmacs;

**editors** such as XEmacs, a general purpose programming editor and tool along with several add-on applications via their supplied elisp code, and the TeXmacs WYSIWYG mathematical editor and interface to several languages and systems;

**miscellaneous** scientific programs such as aplus, aribas, ent, euler, evolver, felt, freefem, gambit, geg, geomview, ghemical, glp, gmt, gperiodic, ipe, lp-solve, lush, mpb, mpqc, and rasmol.

This unique combination of scientific applications, combined with the “live cdrom” technology from Knoppix enables any standard personal computer to be transformed into a scientific workstation.

## 3 Quantian and openMosix Single-System Image clustering

The initial Quantian [6] presentation conjectured that permitting several machines running Quantian to be combined in a single-system image openMosix-style cluster may be a natural extension. Shortly thereafter, the first releases of ClusterKnoppix [3] became available. By basing Quantian on ClusterKnoppix instead of the standard Knoppix system, Quantian has added a new dimension: openMosix clustering support. The next section discusses ClusterKnoppix.

### 3.1 ClusterKnoppix contributions

ClusterKnoppix [3] extends Knoppix by combining it with an openMosix-enabled kernel. OpenMosix [4] is a Linux kernel extension for single-system image (SSI) clustering. This SSI kernel extension turns a network of ordinary computers into a larger computer for Linux applications that presents itself to the user as a single, more powerful computer rather than a collection of machines.

Once openMosix is enabled and started, nodes in the cluster can start talking to one another while continuously attempting to optimize the resource allocation by migrating processes from ‘busy’ nodes to ‘spare’ nodes in order to split the total computing load evenly across the cluster. The resulting system is approximately linearly scalable in the number of nodes. Moreover, with the openMosix auto discovery mechanism, a new node can be added while the cluster is running. The cluster will automatically begin to use the new resource allowing for dynamic sizing of the cluster as well as some level of fault-tolerance.

Applications do not need to be programmed specifically for openMosix (as opposed to Beowulf systems which require explicit communication protocols such as MPI). Since all openMosix extensions are inside the kernel, every Linux application automatically and transparently benefits from the distributed computing concept of openMosix. The cluster behaves much like a Symmetric Multi-Processor (SMP) computer, but scales to well over a thousand nodes (which could themselves be SMPs for a larger number of total processors).

ClusterKnoppix provides these openMosix features in a framework with the following key benefits:

- includes openMosix terminal server using PXE, DHCP and tftp permitting linux clients to boot via the network (which permits to use nodes without hard disk, cdrom or floppy);
- operates openMosix in autodiscovery mode so that new nodes automatically join the cluster minimising the need for configuration or administration;
- contains cluster management tools such as openMosixview;
- setup such that every node has root access to every other node via ssh using RSA-encrypted keys;
- provides Mosix / Direct File System (MFS/DFSA) support which enables all nodes to see each others files;
- permits the choice for each node to be run as a graphical workstation (lab/demo setup) or as a text-only console conserving system memory.

The most recent version of clusterKnoppix also adds the tyd service from the CHAOS project [8, 9]. CHAOS focuses on adding a virtual private network (VPN) layer on top of open network protocols in order for the openMosix-style cluster to be used securely on public networks. To this end, CHAOS includes the FreeSWAN kernel patches.

### 3.2 Quantian and openMosix

By providing the scientific applications listed in section 2.2 in the framework of an openMosix-enabled “live cdrom”, Quantian offers a synthesis of three distinct domains:

1. Knoppix auto-configuration and ease of installation on virtually any recent desktop or laptop;
2. a large number of analytically-focused applications pre-loaded and configured ready for deployment;
3. openMosix extensions to the Linux kernel permitting approximately linearly-scaled performance increases with additional nodes in a single-system image cluster.

## 4 Application example

To be done. Maybe use Monte Carlo simulation or Bootstrap example.

## 5 Summary

To be done.

## References

- [1] Dirk Eddelbuettel, “Quantian”, <http://dirk.eddelbuettel.com/quantian.html>, (2003).
- [2] Klaus Knopper, “Knoppix”, <http://www.knopper.net/knoppix/index-en.html>, (2003).
- [3] Wim Vandersmissen, “ClusterKnoppix”, <http://bofh.be/clusterknoppix>, (2003).
- [4] Moshe Bar et al., “openMosix”, <http://www.openmosix.org>, (2003).
- [5] R Core, “R”, <http://www.r-project.org>, (2003).
- [6] Dirk Eddelbuettel, “Quantian: A Scientific Computing Environment”, Proceedings of the 3rd International Workshop on Distributed Statistical Computing (DSC 2003), ISSN 1609-395X.
- [7] The Comprehensive R Archive Network, <http://cran.r-project.org>, (2003).
- [8] Ian Latter, “CHAOS”, <http://itsecurity.mq.edu.au/chaos>, (2003).
- [9] Ian Latter, “Security and openMosix: Securely deploying SSI cluster technology over untrusted networking infrastructure”, white paper, Macquarie University, (2003).